# NOvA Control Room Computing

Operational Description and Support

Andrew Norman, John Urish
*Fermilab, Computing Division*

CD-DOCDB-4255-v1

(Last Revised on 2/23/11)

## Overview

The NOvA Control room is currently located on the 12th floor of Wilson hall and is in a shared space with other Intensity Frontier (IF) experiments. The control room serves as central point from which operations of the NOvA detectors are conducted and from which the Data Acquisition (DAQ) systems for the detectors are interacted with.

The NOvA DAQ computing model has been designed with the remote nature of the detector systems in mind. The computing resources that acquire data from the NOvA detectors, are housed in a controlled facility that is local to each detector. In the case of the current NOvA near detector, the DAQ computing is located in the Lattice Computing Center (LCC). A high speed (10Gig) fiber trunk line connects LCC to the network in the detector building. For the far detector, the computing will be housed in a dedicated computing facility in the Ash River building, and connected to the detector floor with a set of direct high speed fiber trunk lines that connect to the network switches for each detector di-block area. In both cases the network (DAQNet) is an entirely private network and is firewalled off from external traffic, both for both bandwidth and security concerns.

Interfacing with the DAQ at each site is done through a system of gateway machines, which have access to both the internal DAQNet and the external public network for controls. All of the DAQ software has been designed with this access model in mind, and as a result all the DAQ and system monitoring software can been run transparently from remote sites without any special configuration.

The Wilson Hall control room uses this access model, by providing graphics terminals for shift operators to login to and then connect to the detector gateway machines. The control room machines are used essentially as "X-terminals" and require no specialized hardware, software or system libraries be present to fulfill their rolls.

## System Users

The control room machines use a single experiment specific user access model. Each of the control room machines has a single user account which is used by the experiment operators. Remote access to the this account is handled by inclusion of appropriate users in the .k5login file. This file is maintained by the experiment and is distributed to each of the control machines.

The home directory of the user account is not shared across machines. This decision was made to prevent a single point of failure (a shared disk) from affecting all the

machines. Each home directory contains only a limited set of scripts which are used to start and stop DAQ services. These are maintained from the main DAQ cluster, and can be deployed or restored [on demand] to any machine that is being used to access the DAQ cluster.

All other files stored in the home areas on the control room machines are considered, temporary or non-persistent and do not need to be backed up, archived or in any other fashion preserved.

No other user accounts are maintained on the control room machines.

## Hardware Configuration

The NOvA control room was outfitted for the initial Near Detector on the Surface running, with five Dell brand personal computers. The computers were purchased off of the Fermilab standard desktop computing lists. The machines are Dell OptiPlex model 780s, where each machine has a total of 8GB of system memory and an Intel Core2 Duo cpu, running at 3.3GHz. The only deviation of the control room machines from the base Scientific Linux Fermi configuration is that each machines features a NVIDIA graphics cards with support for quad dvi displays.

Two of the control room machines take advantage of their four display capability and a run in a quad head configuration to provide the primary workspaces for the NOvA DAQ run control interfaces, and for the data quality and monitoring interfaces. Two of the remaining machines are configured as dual head workstations and run a combination of applications that include cluster monitoring and remote video monitoring of the detector building. The remaining machine in the control room is configured as a single head machine and is denoted as a hot spare for control room, and is used as a general purpose terminal for shifters, developers and DAQ experts to perform work on while in the control room.

The DAQ machines are connected to the standard Fermilab public network. The machines are connected to the network via an 8 port switch and fiber connection that are dedicated to the NOvA experiment (i.e. not shared with the other experiments in the same room.) The machines are located on the 131.225.52.x subnet, and have ACL restrictions placed on them to prevent direct access from machines outside of the Fermilab network.

## Software Configuration

The control room machines were designed to be as generic and and interchangeable as possible with regards to both their hardware and software configurations. As a result, each machine has been identically configured with a Scientific Linux Fermi

distribution. The original installation was done using SLF 5.5 with the NOvA DAQ workgroup package set. This workgroup setup was developed by A.Norman, D. Perevolov and others in the NOvA DAQ group to match the base needs of the NOvA DAQ software. The setup involves a set of system packages to provide normal interactive desktop X session functionality. External packages that would normally be required for running experiment specific code, are confined to an experiment specific UPS repository which properly versioned, stabled and frozen libraries and packages.

## Uptime Requirements

The NOvA control room is staffed 24 hours per day, 7 days per week when beam is running to the experiment. The experiment maintains three distinct shift periods running from 07:45-16:15, 15:45-00:15 and 23:45-08:15. Each shift period is staffed by one individual who has been trained in the basic operations of the data acquisition system. This individual is not necessarily an expert with respect to the computing systems but does have limited knowledge of how to start, stop and restart the systems. When the beam is not running to the detector, the control room is staffed at a reduced level, and most of the work taking place pertains to calibration and studies of the detector.

During normal operations, *no one control room computer is considered critical to the operation of the facility*. Each machine's hardware and software configuration was designed to be generic enough that in the event of a system failure, any other machine in the control room could duplicate the functionality of any other. Due to the number of software displays and DAQ components that must be visible at any one time to the shift operator, a minmum of 10 screens spread across three distinct machines are required for operation of the NOvA DAQ and monitoring. Under this mode of operation, the control room can currently experience at most 2 simultaneous full system failures at one time without out suffering any loss of controls functionality.

Due to this ability to run with only 60% of the systems operational, we foresee 8x5 support to be sufficient for standard maintenance. Catastrophic events that would affect or cause failure in more than 60% of the control room machines, would require 24x7 support during beam operations, and 8x5 support during periods when the beam is not running.

We would also request that maintenance that requires system down time be scheduled in advance, and if possible scheduled during known beam down times. When maintenance need must unavoidably coincide with data taking, we would request that the maintenance be phased in such a manner as to take down no more than 40% of our control room machines at one time, thereby allowing operations and data taking to continue with minimal disruption.

# Appendix I – Machine Configurations

Nova-116121 June 9, 2010 Nova-daq-01.fnal.gov Configuration

Dell Optiplex 780
Service Tag CJJJ7M1
CPU = Core 2 Duo E8600 3.33 GHz
RAM = 8GB
video = NVIDIA G98 (Quadro NVS 420)
Sound = Intel onboard 82801 ICHIO Family

```
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/mapper/isw_fiihacbd_ARRAYp1
                      99188468   4062640  90005916   5% /
/dev/mapper/isw_fiihacbd_ARRAYp6
                     167870844    192012 159013776   1% /data
/dev/mapper/isw_fiihacbd_ARRAYp5
                       9920592    155052   9253472   2% /home
/dev/mapper/isw_fiihacbd_ARRAYp2
                       9920624    251364   9157192   3% /var
```

Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       12748   102398278+  83  Linux
/dev/sda2           12749       14023    10241437+  83  Linux
/dev/sda3           14024       16063    16386300   82  Linux swap / Solaris
/dev/sda4           16064       38913   183542625    5  Extended
/dev/sda5           16064       17338    10241406   83  Linux
/dev/sda6           17339       38913   173301156   83  Linux
```

Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

```
   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *           1       12748   102398278+  83  Linux
/dev/sdb2           12749       14023    10241437+  83  Linux
/dev/sdb3           14024       16063    16386300   82  Linux swap / Solaris
/dev/sdb4           16064       38913   183542625    5  Extended
/dev/sdb5           16064       17338    10241406   83  Linux
/dev/sdb6           17339       38913   173301156   83  Linux
```

NIC = Intel onboard 82567LM-3 Gigabit Ethernet (Internal NIC)
MAC = 00-25-64-c9-8b-82

IP = 131.225.52.118
Subnet Mask = 255.255.252.0
Default Gateway = 131.225.55.200

NFS auto mounts:

Software:
Scientific Linux Fermi LTS release 5.5
Kernel 2.6.18-194.3.1.el5PAE #1 SMP
Nova DAQ Workgroup = 9553

Nova-116123 June 9, 2010 Nova-daq-02.fnal.gov Configuration

Dell Optiplex 780
Service Tag 6JJJ7M1

CPU = Core 2 Duo E8600 3.33 GHz
RAM = 8GB
video = NVIDIA G98 (Quadro NVS 420)
Sound = Intel onboard 82801 ICHIO Family

```
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/mapper/isw_bdgfaeagfb_ARRAYp1
                      99188468   4062612  90005944   5% /
/dev/mapper/isw_bdgfaeagfb_ARRAYp6
                     167870844    192012 159013776   1% /data
/dev/mapper/isw_bdgfaeagfb_ARRAYp3
                       9920624    282924   9125632   4% /var
/dev/mapper/isw_bdgfaeagfb_ARRAYp5
                       9920592    156216   9252308   2% /home
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       12748   102398278+  83  Linux
/dev/sda2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sda3           14789       16063    10241437+  83  Linux
/dev/sda4           16064       38913   183542625    5  Extended
/dev/sda5           16064       17338    10241406   83  Linux
/dev/sda6           17339       38913   173301156   83  Linux

Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *           1       12748   102398278+  83  Linux
/dev/sdb2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sdb3           14789       16063    10241437+  83  Linux
/dev/sdb4           16064       38913   183542625    5  Extended
/dev/sdb5           16064       17338    10241406   83  Linux
/dev/sdb6           17339       38913   173301156   83  Linux
```

NIC = Intel onboard 82567LM-3 Gigabit Ethernet (Internal NIC)
MAC = 00-25-64-c9-88-c7
IP = 131.225.52.134
Subnet Mask = 255.255.252.0
Default Gateway = 131.225.55.200

NFS auto mounts:

Software:
Scientific Linux Fermi LTS release 5.5
Kernel 2.6.18-194.3.1.el5PAE #1 SMP
Nova DAQ Workgroup = 9553

Nova-116120 December 15, 2010 Nova-daq-03.fnal.gov Configuration

Dell Optiplex 780
Service Tag 6HJJ7M1
CPU = Core 2 Duo E8600 3.33 GHz
RAM = 8GB
video = NVIDIA G98 (Quadro NVS 420)
Sound = Intel onboard 82801 ICHIO Family

```
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/mapper/isw_geebajee_ARRAYp1
                      99188468   4062268  90006288   5% /
/dev/mapper/isw_geebajee_ARRAYp6
                     167870844    192012 159013776   1% /data
/dev/mapper/isw_geebajee_ARRAYp3
                       9920624    251480   9157076   3% /var
/dev/mapper/isw_geebajee_ARRAYp5
                       9920592    155052   9253472   2% /home
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       12748   102398278+  83  Linux
/dev/sda2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sda3           14789       16063    10241437+  83  Linux
/dev/sda4           16064       38913   183542625    5  Extended
/dev/sda5           16064       17338    10241406   83  Linux
/dev/sda6           17339       38913   173301156   83  Linux

Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *           1       12748   102398278+  83  Linux
/dev/sdb2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sdb3           14789       16063    10241437+  83  Linux
/dev/sdb4           16064       38913   183542625    5  Extended
/dev/sdb5           16064       17338    10241406   83  Linux
/dev/sdb6           17339       38913   173301156   83  Linux
```

NIC = Intel onboard 82567LM-3 Gigabit Ethernet (Internal NIC)
MAC = 00-25-64-c9-88-65
IP = 131.225.53.34
Subnet Mask = 255.255.252.0
Default Gateway = 131.225.55.200

NFS auto mounts:

Software:
Scientific Linux Fermi LTS release 5.5
Kernel 2.6.18-194.3.1.elPAE #1 SMP
Nova DAQ Workgroup

Nova-116119 June 9, 2010 Nova-daq-04.fnal.gov Configuration

Dell Optiplex 780
Service Tag GHJJ7M1
CPU = Core 2 Duo E8600 3.33 GHz
RAM = 8GB
video = NVIDIA G98 (Quadro NVS 420)
Sound = Intel onboard 82801 ICHIO Family

```
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/mapper/isw_cebicdbjeh_ARRAYp1
                      99188468   4062256  90006300   5% /
/dev/mapper/isw_cebicdbjeh_ARRAYp6
                     167870844    192012 159013776   1% /data
/dev/mapper/isw_cebicdbjeh_ARRAYp3
                       9920624    252892   9155664   3% /var
/dev/mapper/isw_cebicdbjeh_ARRAYp5
                       9920592    156588   9251936   2% /home
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       12748   102398278+  83  Linux
/dev/sda2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sda3           14789       16063    10241437+  83  Linux
/dev/sda4           16064       38913   183542625    5  Extended
/dev/sda5           16064       17338    10241406   83  Linux
/dev/sda6           17339       38913   173301156   83  Linux

Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *           1       12748   102398278+  83  Linux
/dev/sdb2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sdb3           14789       16063    10241437+  83  Linux
/dev/sdb4           16064       38913   183542625    5  Extended
/dev/sdb5           16064       17338    10241406   83  Linux
/dev/sdb6           17339       38913   173301156   83  Linux
```

NIC = Intel onboard 82567LM-3 Gigabit Ethernet (Internal NIC)
MAC = 00-25-64-c9-86-b3
IP = 131.225.53.184
Subnet Mask = 255.255.252.0
Default Gateway = 131.225.55.200

NFS auto mounts:

Software:
Scientific Linux Fermi LTS release 5.5
Kernel 2.6.18-194.3.1.el5PAE #1 SMP
Nova DAQ Workgroup = 9553

Nova-116122 June 9, 2010 Nova-shifter-1.fnal.gov Configuration

Dell Optiplex 780
Service Tag 6JJJ7M1
CPU = Core 2 Duo E8600 3.33 GHz
RAM = 8GB
video = NVIDIA G98 (Quadro NVS 420)
Sound = Intel onboard 82801 ICHIO Family

```
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/mapper/isw_bjfdadhjga_ARRAYp1
                      99188468   4062256  90006300   5% /
/dev/mapper/isw_bjfdadhjga_ARRAYp6
                     167870844    192012 159013776   1% /data
/dev/mapper/isw_bjfdadhjga_ARRAYp3
                       9920624    252856   9155700   3% /var
/dev/mapper/isw_bjfdadhjga_ARRAYp5
                       9920592    156204   9252320   2% /home
Disk /dev/sda: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1   *           1       12748   102398278+  83  Linux
/dev/sda2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sda3           14789       16063    10241437+  83  Linux
/dev/sda4           16064       38913   183542625    5  Extended
/dev/sda5           16064       17338    10241406   83  Linux
/dev/sda6           17339       38913   173301156   83  Linux

Disk /dev/sdb: 320.0 GB, 320072933376 bytes
255 heads, 63 sectors/track, 38913 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1   *           1       12748   102398278+  83  Linux
/dev/sdb2           12749       14788    16386300   82  Linux swap / Solaris
/dev/sdb3           14789       16063    10241437+  83  Linux
/dev/sdb4           16064       38913   183542625    5  Extended
/dev/sdb5           16064       17338    10241406   83  Linux
/dev/sdb6           17339       38913   173301156   83  Linux
```

NIC = Intel onboard 82567LM-3 Gigabit Ethernet (Internal NIC)
MAC = 00-25-64-c9-88-80
IP = 131.225.53.154
Subnet Mask = 255.255.252.0
Default Gateway = 131.225.55.200

NFS auto mounts:

Software:
Scientific Linux Fermi LTS release 5.5
Kernel 2.6.18-194.3.1.el5PAE #1 SMP
Nova DAQ Workgroup = 9553